Combining Contextuality and Causality

Samson Abramsky, Rui Soares Barbosa and Amy Searle

Department of Computer Science, UCL
Iberian Nanotechnology Laboratory, Braga
Department of Physics, University of Oxford

QCQMB Prague December 2022

Combining contextuality and causality

Reasons:

Combining contextuality and causality

Reasons:

- Causality is a significant additional structure, which may arise from:
  - causal structure of an experiment
  - feed-forward in MBQC; adaptive computation
  - at a fundamental level: the causal structure of spacetime

Combining contextuality and causality

Reasons:

- Causality is a significant additional structure, which may arise from:
  - causal structure of an experiment
  - feed-forward in MBQC; adaptive computation
  - at a fundamental level: the causal structure of spacetime
- Causal inference, once shunned by statisticians, has become very influential across a wide range of disciplines, with seminal work by Pearl *et al.*

Combining contextuality and causality

Reasons:

- Causality is a significant additional structure, which may arise from:
  - causal structure of an experiment
  - feed-forward in MBQC; adaptive computation
  - at a fundamental level: the causal structure of spacetime

- Causal inference, once shunned by statisticians, has become very influential across a wide range of disciplines, with seminal work by Pearl *et al.*

Our objectives:

Combining contextuality and causality

Reasons:

- Causality is a significant additional structure, which may arise from:
  - causal structure of an experiment
  - feed-forward in MBQC; adaptive computation
  - at a fundamental level: the causal structure of spacetime
- Causal inference, once shunned by statisticians, has become very influential across a wide range of disciplines, with seminal work by Pearl *et al.*

Our objectives:

- A more fine-grained analysis of contextuality: no-signalling/no-disturbance **outside the causal past** (the light cone)

Combining contextuality and causality

Reasons:

- Causality is a significant additional structure, which may arise from:
  - causal structure of an experiment
  - feed-forward in MBQC; adaptive computation
  - at a fundamental level: the causal structure of spacetime
- Causal inference, once shunned by statisticians, has become very influential across a wide range of disciplines, with seminal work by Pearl *et al.*

Our objectives:

- A more fine-grained analysis of contextuality: no-signalling/no-disturbance **outside the causal past** (the light cone)
- Better connection with computational models such as circuits and MBQC, deeper analysis of contextuality and quantum advantage

Causal networks

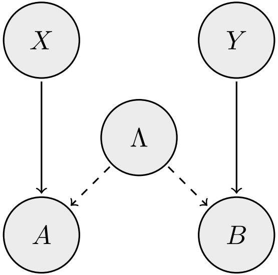

Bell scenario

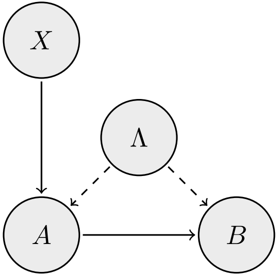

Instrumental scenario

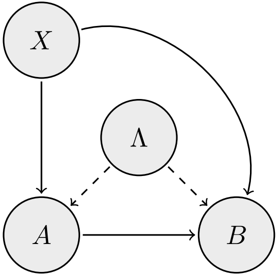

Variant

Causal networks

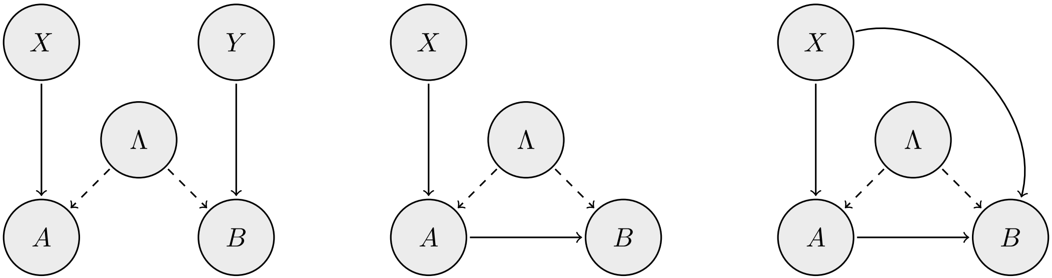

Bell scenario Instrumental scenario Variant

Note that X and Y are random variables whose values are the **measurement settings**; A and B are the **measurement outcomes**; Λ is a **latent variable**.

Our approach

We seek to extend the Abramsky-Brandenburger sheaf-theoretic approach to contextuality to incorporate causal structure.

Our approach

We seek to extend the Abramsky-Brandenburger sheaf-theoretic approach to contextuality to incorporate causal structure.

This is a general, mathematically robust approach, which provides a basis for:

- the contextual fraction
- general characterisation of contextual inequalities in terms of consistency conditions ("logical Bell inequalities", Boole's "conditions of possible experience")
- resource theory of contextuality
- simulations between contextual systems
- cohomological criteria for contextuality; topology of contextuality
- connections with logic and computation; database theory, constraint satisfaction
- generalized Vorob'ev theorem

Our approach

We seek to extend the Abramsky-Brandenburger sheaf-theoretic approach to contextuality to incorporate causal structure.

This is a general, mathematically robust approach, which provides a basis for:

- the contextual fraction
- general characterisation of contextual inequalities in terms of consistency conditions ("logical Bell inequalities", Boole's "conditions of possible experience")
- resource theory of contextuality
- simulations between contextual systems
- cohomological criteria for contextuality; topology of contextuality
- connections with logic and computation; database theory, constraint satisfaction
- generalized Vorob'ev theorem

The hope is that these features will all carry over to the refined version incorporating causality.

Comparison with CbD

- Contextuality-by-Default (CbD), Dzhafarov, Kujala *et al.*

Comparison with CbD

- Contextuality-by-Default (CbD), Dzhafarov, Kujala *et al.*

- In CbD, **every** variable is regarded as contextual, differently labelled in each context.
  One then measures the extent to which different occurrences with the same "content" can be regarded as the same.
  This allows for the analysis of arbitrary signalling systems.

Comparison with CbD

- Contextuality-by-Default (CbD), Dzhafarov, Kujala *et al.*

- In CbD, **every** variable is regarded as contextual, differently labelled in each context. One then measures the extent to which different occurrences with the same "content" can be regarded as the same.
  This allows for the analysis of arbitrary signalling systems.

- By contrast, we wish to **explicitly describe** a given causal background, which might arise from the structure of an experiment, circuit, or physical system.

Comparison with CbD

- Contextuality-by-Default (CbD), Dzhafarov, Kujala *et al.*

- In CbD, **every** variable is regarded as contextual, differently labelled in each context. One then measures the extent to which different occurrences with the same "content" can be regarded as the same.
  This allows for the analysis of arbitrary signalling systems.

- By contrast, we wish to **explicitly describe** a given causal background, which might arise from the structure of an experiment, circuit, or physical system.

- Signalling is then allowed within the backwards light-cone or causal past of an event, while no-signalling is required outside it.

Comparison with CbD

- Contextuality-by-Default (CbD), Dzhafarov, Kujala *et al.*

- In CbD, **every** variable is regarded as contextual, differently labelled in each context. One then measures the extent to which different occurrences with the same "content" can be regarded as the same.
  This allows for the analysis of arbitrary signalling systems.

- By contrast, we wish to **explicitly describe** a given causal background, which might arise from the structure of an experiment, circuit, or physical system.

- Signalling is then allowed within the backwards light-cone or causal past of an event, while no-signalling is required outside it.

- The aim is to recognise an extra level of structure which is relevant for our purposes.

Precursors

- Pearl had already noted the connection with Bell inequalities in his seminal paper on testability of causal models with latent and instrumental variables
  J. Pearl, *Proc. Conf. Uncertainty in AI*, 1995.

- The extension of causal networks to allow for quantum resources, or more generally the operations offered by Generalized Probabilistic Theories, studied e.g. in:
  - Henson, Lal and Pusey, *NJP*, 2014.
  - Chaves, Garvacho, Agresti, Giulio, Aolita, Giacomini and Sciarrino, *Nature Physics*, 2018.
  - van Himbeek, Brask, Pironio, Ramanathan, Sainz and Wolfe, *Quantum*, 2019.

  .

- Shane Mansfield (lectures 2016) studied a refinement of the sheaf-theoretic approach with an order on the measurements, and used this to study the two-slit experiment and Leggett-Garg.

- Stefano Gogioso and Nicola Pinzani (QPL 2021, 474(!)-page arXiv paper 2022) study a causal refinement of the sheaf-theoretic approach for **Bell scenarios**.

Brief recap of the sheaf-theoretic approach

*SA, Brandenburger, *New Journal of Physics*, 2011.

in\out	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a,b)	–	–	–	–
(a,b')	–	–	–	–
(a',b)	–	–	–	–
(a',b')	–	–	–	–

A **measurement scenario** $\mathbf{X} = \langle X, \Sigma, O \rangle$:

- X – a finite set of measurements
- Σ – a simplicial complex on X
  faces are called the **measurement contexts**
- $O = (O_x)_{x \in X}$ – for each $x \in X$ a finite non-empty set of possible outcomes O_x

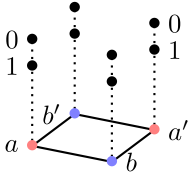

Brief recap of the sheaf-theoretic approach

*SA, Brandenburger, *New Journal of Physics*, 2011.

A **measurement scenario** $\mathbf{X} = \langle X, \Sigma, O \rangle$:

- X – a finite set of measurements
- Σ – a simplicial complex on X
  faces are called the **measurement contexts**
- $O = (O_x)_{x \in X}$ – for each $x \in X$ a finite non-empty set of possible outcomes O_x

in\out	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a, b)	$1/2$	0	0	$1/2$
(a, b')	$1/2$	0	0	$1/2$
(a', b)	$1/2$	0	0	$1/2$
(a', b')	0	$1/2$	$1/2$	0

An **empirical model** $e = \{e_\sigma\}_{e \in \Sigma}$ on $\mathbf{X}$:

- Each e_σ is a prob. distribution over joint outcomes $\prod_{x \in \sigma} O_x$ for σ
- **generalised no-signalling** holds:
  $\forall \sigma, \tau \in \Sigma, \sigma \subseteq \tau.$

$$e_\tau|_\sigma = e_\sigma$$

(i.e. marginals are well-defined)

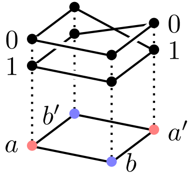

Brief recap of the sheaf-theoretic approach

*SA, Brandenburger, *New Journal of Physics*, 2011.

A **measurement scenario** $\mathbf{X} = \langle X, \Sigma, O \rangle$:

- X – a finite set of measurements
- Σ – a simplicial complex on X
  faces are called the **measurement contexts**
- $O = (O_x)_{x \in X}$ – for each $x \in X$ a finite non-empty set of possible outcomes O_x

in\out	$(0, 0)$	$(0, 1)$	$(1, 0)$	$(1, 1)$
(a, b)	$1/2$	0	0	$1/2$
(a, b')	$1/2$	0	0	$1/2$
(a', b)	$1/2$	0	0	$1/2$
(a', b')	0	$1/2$	$1/2$	0

An **empirical model** $e = \{e_\sigma\}_{e \in \Sigma}$ on $\mathbf{X}$:

- Each e_σ is a prob. distribution over joint outcomes $\prod_{x \in \sigma} O_x$ for σ
- **generalised no-signalling** holds:
  $\forall \sigma, \tau \in \Sigma, \sigma \subseteq \tau$.

$$e_\tau|_\sigma = e_\sigma$$

(i.e. marginals are well-defined)

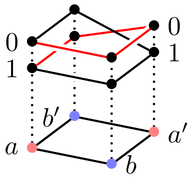

Desiderata

Our aim is to carry this structure over to the sheaf-theoretic setting for contextuality, yielding the desired causal refinement.

Desiderata

Our aim is to carry this structure over to the sheaf-theoretic setting for contextuality, yielding the desired causal refinement.

The hope is that this leads to a smooth extension of the theory, to which all the current aspects:

contextual fraction, logical Bell inequalities, resource theory, simulations,
cohomological criteria, connections with logic and computation, etc. etc.

also lift smoothly.

Desiderata

Our aim is to carry this structure over to the sheaf-theoretic setting for contextuality, yielding the desired causal refinement.

The hope is that this leads to a smooth extension of the theory, to which all the current aspects:

contextual fraction, logical Bell inequalities, resource theory, simulations, cohomological criteria, connections with logic and computation, etc. etc.

also lift smoothly.

Grades of causal involvement:

- global ordering on measurements (cf. Mansfield and Gogioso-Pinzani approaches)
- dependence on measurement outcomes, allowing e.g. for feed-forward in MBQC, adaptive computation
- recognizing the different roles played by Nature and Experimenter in their interactions

Dual faces of causality

Causality may be:

- imposed by Nature – a **causal background**
- imposed by the experimenter, e.g. to achieve computational effects (adaptive computation).

Dual faces of causality

Causality may be:

- imposed by Nature – a **causal background**
- imposed by the experimenter, e.g. to achieve computational effects (adaptive computation).

We will illustrate these two sources of causality in two basic examples.

Example I: causal background a la G-P

Standard Bell-CHSH bipartite scenario: Alice and Bob, with local measurements M_A and M_B, and outcomes O_A and O_B.

Example I: causal background a la G-P

Standard Bell-CHSH bipartite scenario: Alice and Bob, with local measurements M_A and M_B, and outcomes O_A and O_B.

We assume that Alice's events **causally precede** those of Bob.

Example I: causal background a la G-P

Standard Bell-CHSH bipartite scenario: Alice and Bob, with local measurements M_A and M_B, and outcomes O_A and O_B.

We assume that Alice's events **causally precede** those of Bob.

Thus Bob's backward light-cone includes the events where Alice chooses a measurement and observes an outcome.

Example I: causal background a la G-P

Standard Bell-CHSH bipartite scenario: Alice and Bob, with local measurements M_A and M_B, and outcomes O_A and O_B.

We assume that Alice's events **causally precede** those of Bob.

Thus Bob's backward light-cone includes the events where Alice chooses a measurement and observes an outcome.

Whereas in a standard, "flat" scenario, we would have deterministic outcomes given by functions

$$s_A : M_A \longrightarrow O_A, \quad s_B : M_B \longrightarrow O_B,$$

with these causal constraints, we have functions

$$s_A : M_A \longrightarrow O_A, \quad s_B : M_A \times M_B \longrightarrow O_B$$

Example I: causal background a la G-P

Standard Bell-CHSH bipartite scenario: Alice and Bob, with local measurements M_A and M_B, and outcomes O_A and O_B.

We assume that Alice's events **causally precede** those of Bob.

Thus Bob's backward light-cone includes the events where Alice chooses a measurement and observes an outcome.

Whereas in a standard, "flat" scenario, we would have deterministic outcomes given by functions

$$s_A : M_A \longrightarrow O_A, \quad s_B : M_B \longrightarrow O_B,$$

with these causal constraints, we have functions

$$s_A : M_A \longrightarrow O_A, \quad s_B : M_A \times M_B \longrightarrow O_B$$

That is, the responses by Nature to Bob's measurement may depend on the previous measurement made by Alice.

Example I ctd

If we have measurements $x_1, x_2 \in M_A$, $y \in M_B$, then we can have $\{(x_1, 0), (y, 0)\}$ and $\{(x_2, 0), (y, 1)\}$ as valid histories in a single deterministic model.

Example I ctd

If we have measurements $x_1, x_2 \in M_A$, $y \in M_B$, then we can have $\{(x_1, 0), (y, 0)\}$ and $\{(x_2, 0), (y, 1)\}$ as valid histories in a single deterministic model.

Of the usual no-signalling/compatibility equations

$$\begin{array}{llcl} (1) & e_{\{x_i, y\}}\big|_{\{x_i\}} & = & e_{\{x_i\}} \\ (2) & e_{\{x_i, y\}}\big|_{\{y\}} & = & e_{\{y\}} \end{array}$$

only (1) remains: $e_{\{y\}}$ is not even defined, since $\{y\}$ is not a "causally secured" context.

Example I ctd

If we have measurements $x_1, x_2 \in M_A$, $y \in M_B$, then we can have $\{(x_1, 0), (y, 0)\}$ and $\{(x_2, 0), (y, 1)\}$ as valid histories in a single deterministic model.

Of the usual no-signalling/compatibility equations

$$
\begin{array}{llcl}
(1) & e_{\{x_i, y\}}|_{\{x_i\}} & = & e_{\{x_i\}} \\
(2) & e_{\{x_i, y\}}|_{\{y\}} & = & e_{\{y\}}
\end{array}
$$

only (1) remains: $e_{\{y\}}$ is not even defined, since $\{y\}$ is not a "causally secured" context.

Thus no-signalling is relaxed in a controlled fashion.

Example II: Anders-Browne (PRL 2009)

This shows how we can use a form of Experimenter-imposed causality to promote two sub-universal computational models (Pauli measurements and mod-2 linear classical processing) to universal MBQC.

Example II: Anders-Browne (PRL 2009)

This shows how we can use a form of Experimenter-imposed causality to promote two sub-universal computational models (Pauli measurements and mod-2 linear classical processing) to universal MBQC.

Uses GHZ state as a resource state: GHZ $= \frac{|\uparrow\uparrow\uparrow\rangle + |\downarrow\downarrow\downarrow\rangle}{\sqrt{2}}$.

Example II: Anders-Browne (PRL 2009)

This shows how we can use a form of Experimenter-imposed causality to promote two sub-universal computational models (Pauli measurements and mod-2 linear classical processing) to universal MBQC.

Uses GHZ state as a resource state: GHZ $= \frac{|\uparrow\uparrow\uparrow\rangle + |\downarrow\downarrow\downarrow\rangle}{\sqrt{2}}$.

	$+++$	$++-$	$+-+$	$+--$	$-++$	$-+-$	$--+$	$---$
XYY	0	1	1	0	1	0	0	1
YXY	0	1	1	0	1	0	0	1
YYX	0	1	1	0	1	0	0	1
XXX	1	0	0	1	0	1	1	0

Example II: Anders-Browne (PRL 2009)

This shows how we can use a form of Experimenter-imposed causality to promote two sub-universal computational models (Pauli measurements and mod-2 linear classical processing) to universal MBQC.

Uses GHZ state as a resource state: GHZ $= \frac{|\uparrow\uparrow\uparrow\rangle + |\downarrow\downarrow\downarrow\rangle}{\sqrt{2}}$.

	$+++$	$++-$	$+-+$	$+--$	$-++$	$-+-$	$--+$	$---$
XYY	0	1	1	0	1	0	0	1
YXY	0	1	1	0	1	0	0	1
YYX	0	1	1	0	1	0	0	1
XXX	1	0	0	1	0	1	1	0

In terms of parities (product of $+1/-1$ outputs):

$$\begin{array}{ccccc} X_1 & Y_2 & Y_3 & = & -1 \\ Y_1 & X_2 & Y_3 & = & -1 \\ Y_1 & Y_2 & X_3 & = & -1 \\ X_1 & X_2 & X_3 & = & +1 \end{array}$$

Using Experimenter causal flow to implement AND

Taking X as 0, Y as 1, we consider the measurements for Alice and Bob as inputs to an AND-gate.

Using Experimenter causal flow to implement AND

Taking X as 0, Y as 1, we consider the measurements for Alice and Bob as inputs to an AND-gate.

We then use the following simple mapping (XOR on the bit representations) from the Alice-Bob measurements to Charlie's measurement to get the AND-function, which we can read off from the XOR of the outcome bits:

$$\begin{array}{rcl} 0,1 & \mapsto & 1 \\ 1,0 & \mapsto & 1 \\ 1,1 & \mapsto & 0 \\ 0,0 & \mapsto & 0 \end{array} \qquad\qquad \begin{array}{rcl} X,Y & \mapsto & Y \\ Y,X & \mapsto & Y \\ Y,Y & \mapsto & X \\ X,X & \mapsto & X \end{array}$$

Using Experimenter causal flow to implement AND

Taking X as 0, Y as 1, we consider the measurements for Alice and Bob as inputs to an AND-gate.

We then use the following simple mapping (XOR on the bit representations) from the Alice-Bob measurements to Charlie's measurement to get the AND-function, which we can read off from the XOR of the outcome bits:

$$\begin{array}{rcl} 0,1 & \mapsto & 1 \\ 1,0 & \mapsto & 1 \\ 1,1 & \mapsto & 0 \\ 0,0 & \mapsto & 0 \end{array} \qquad\qquad \begin{array}{rcl} X,Y & \mapsto & Y \\ Y,X & \mapsto & Y \\ Y,Y & \mapsto & X \\ X,X & \mapsto & X \end{array}$$

Note that:

- this is purely causality employed by the Experimenter; from Nature's point of view, it is the standard GHZ construction
- the above is a simplified "one-shot" description; really there is a (classically computed) feed-forward of measurement settings needed to represent circuits with embedded AND-gates

Game semantics of causality

We shall conceptualise the dual nature of causality as a two-person game, played between Experimenter and Nature:

- The Experimenter's moves are the choices of measurements to be performed.
- Nature's moves are the outcomes.

Game semantics of causality

We shall conceptualise the dual nature of causality as a two-person game, played between Experimenter and Nature:

- The Experimenter's moves are the choices of measurements to be performed.
- Nature's moves are the outcomes.

By formalising this, we can develop a theory of causal contextuality which recovers:

- the usual, "flat" contextuality
- the G-P theory of non-locality in a causal background
- MBQC with adaptive computation
- classical causal networks

as special cases, and more.

Game semantics of causality

We shall conceptualise the dual nature of causality as a two-person game, played between Experimenter and Nature:

- The Experimenter's moves are the choices of measurements to be performed.
- Nature's moves are the outcomes.

By formalising this, we can develop a theory of causal contextuality which recovers:

- the usual, "flat" contextuality
- the G-P theory of non-locality in a causal background
- MBQC with adaptive computation
- classical causal networks

as special cases, and more.

Note Our formalisation will use ideas from Computer Science: Kahn-Plotkin concrete domains and their representations.

Contextuality scenarios

A (flat) contextuality scenario is $(X, O, \mathcal{C})$, where:

- X is a set of **measurements**.
- $O = \{O_x\}_{x \in X}$ is the set of possible **outcomes** for each measurement.
- $\mathcal{C}$ is a **cover**, i.e. a family $\{C_i\}_{i \in I}$ of subsets $C_i \subseteq X$ such that $\bigcup_{i \in I} C_i = X$.

An **event** has the form (x, o), where $x \in X$ and $o \in O_x$. It corresponds to the measurement x being performed, with outcome o.
Given a set of events s,

$$\mathrm{dom}(s) := \pi_1 S = \{x \mid \exists o.\, (x, o) \in s\}.$$

We say that s is **consistent** if

1. for some $C \in \mathcal{C}$, $\mathrm{dom}(s) \subseteq C$;
2. $(x, y), (x, y') \in s$ implies $y = y'$.

In this case, s defines a function from the measurements in its domain to outcomes.
A consistent set of events is a **section**.

Causal contextuality scenarios

A causal contextuality scenario is $(X, O, \mathcal{C}, \vdash)$, where the additional ingredient is an **enabling relation**, which expresses causal constraints.

Causal contextuality scenarios

A causal contextuality scenario is $(X, O, \mathcal{C}, \vdash)$, where the additional ingredient is an **enabling relation**, which expresses causal constraints.

The intended interpretation of $s \vdash x$, where s is a section and $x \in X$, is that it is possible to perform x after the events in s have occurred.

Causal contextuality scenarios

A causal contextuality scenario is $(X, O, \mathcal{C}, \vdash)$, where the additional ingredient is an **enabling relation**, which expresses causal constraints.

The intended interpretation of $s \vdash x$, where s is a section and $x \in X$, is that it is possible to perform x after the events in s have occurred.

Note that this constraint refers to the measurement outcomes as well as the measurements which have been performed. This allows adaptive behaviours to be described.

Histories

Histories

Given such a causal contextuality scenario M, we can use it to generate a set of **histories**, i.e. of sets of events which can happen in a causally consistent fashion. We associate each measurement x with a unique event occurrence, so histories are required to be consistent.

Histories

Given such a causal contextuality scenario M, we can use it to generate a set of **histories**, i.e. of sets of events which can happen in a causally consistent fashion. We associate each measurement x with a unique event occurrence, so histories are required to be consistent.

We define the **accessibility relation** $s \triangleright x$ between sections s and measurements x: $s \triangleright x$ iff $x \notin \text{dom}(s)$, $\text{dom}(s) \cup \{x\} \subseteq C$ for some $C \in \mathcal{C}$, and for some $S \subseteq s$, $S \vdash x$.

Histories

Given such a causal contextuality scenario M, we can use it to generate a set of **histories**, i.e. of sets of events which can happen in a causally consistent fashion. We associate each measurement x with a unique event occurrence, so histories are required to be consistent.

We define the **accessibility relation** $s \triangleright x$ between sections s and measurements x:
$s \triangleright x$ iff $x \notin \mathrm{dom}(s)$, $\mathrm{dom}(s) \cup \{x\} \subseteq C$ for some $C \in \mathcal{C}$, and for some $S \subseteq s$, $S \vdash x$.

Now we can define $\mathcal{H}(M)$, the set of histories over M, inductively by

$$\begin{array}{lcl} H_0 & := & \{\varnothing\} \\ H_{k+1} & := & H_k \;\cup\; \{s \cup \{(x,o)\} \mid s \in H_k, s \triangleright x, o \in O_x\}. \end{array}$$

If X is finite, for some k we will have $H_k = H_{k+1}$, and we take $\mathcal{H}(M) = H_k$ for the least such k.

Example: instrumental scenario

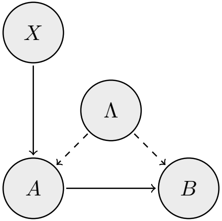

Outcomes: $\{1, 2\}$
Measurement settings

- for Alice: $\{x_1, x_2\}$
- for Bob: $\{y_1, y_2\}$

Enablings:

$$\varnothing \vdash x_i, \qquad (x_i, j) \vdash y_j$$

Thus Alice's measurement outcome determines Bob's measurement setting, without any information as to what Alice's measurement setting was.
The variant where there **is** such information flow can also be represented.

Strategies

We can regard a causal contextuality scenario $M = (X, O, \mathcal{C}, \vdash)$ as specifying a **game** between Experimenter and Nature:

- Events (x, o) correspond to the Experimenter choosing a measurement x, and Nature responding with outcome o.
- The histories correspond to the **plays** or runs of the game.

Strategies

We can regard a causal contextuality scenario $M = (X, O, \mathcal{C}, \vdash)$ as specifying a **game** between Experimenter and Nature:

- Events (x, o) correspond to the Experimenter choosing a measurement x, and Nature responding with outcome o.
- The histories correspond to the **plays** or runs of the game.

Given this interpretation, we define a **strategy for Nature** over the game M as a set of histories $\sigma \subseteq \mathcal{H}(M)$ satisfying the following conditions:

- σ is downwards closed: if $s, t \in \mathcal{H}(M)$ and $s \subseteq t \in \sigma$, then $s \in \sigma$.
- σ is deterministic and total: $\varnothing \in \sigma$, and if $s \in \sigma$ and $s \rhd x$, then there is a unique $o \in O_x$ such that $s \cup \{(x, o)\} \in \sigma$.

Thus in any position s reachable under σ, it has a unique response to any measurement which can be chosen by the Experimenter.

The sheaf of strategies

Given a causal contextuality scenario $M = (X, O, \mathcal{C}, \vdash)$, we can define a presheaf

$$\Gamma : \mathcal{P}(X)^{\mathsf{op}} \longrightarrow \mathbf{Set}$$

For each $U \subseteq X$, $\Gamma(U)$ is the set of strategies for M_U, the restriction of the scenario to measurements in U.

Given $U \subseteq V$, the restriction map $\Gamma(U \subseteq V) : \Gamma(V) \longrightarrow \Gamma(U)$ is given by $\sigma \mapsto \sigma|_U := \sigma \cap \mathcal{H}(M_U)$.

Proposition

Γ *is a presheaf, and satisfies the sheaf condition "above the cover".*

Running the sheaf theory script

We can now follow the same structure as in Abramsky-Brandenburger, replacing the "flat" event sheaf of local sections by the sheaf of strategies.

Running the sheaf theory script

We can now follow the same structure as in Abramsky-Brandenburger, replacing the "flat" event sheaf of local sections by the sheaf of strategies.

We recall the distribution monad $\mathcal{D}_R$, where R is a semiring; when R is the non-negative reals, we recover discrete probability distributions. We have the presheaf $\mathcal{D}_R\Gamma$, obtained by composing the endofunctor part of the monad with Γ.

Running the sheaf theory script

We can now follow the same structure as in Abramsky-Brandenburger, replacing the “flat” event sheaf of local sections by the sheaf of strategies.

We recall the distribution monad $\mathcal{D}_R$, where R is a semiring; when R is the non-negative reals, we recover discrete probability distributions. We have the presheaf $\mathcal{D}_R\Gamma$, obtained by composing the endofunctor part of the monad with Γ.

An **empirical model** over the scenario $(M, \mathcal{C})$ is a family $\{e_i\}$, where $e_i \in \mathcal{D}_R\Gamma(C_i)$, subject to the usual compatibility conditions: for all i, j, $e_i|_{C_i \cap C_j} = e_j|_{C_i \cap C_j}$. Thus e_i assigns a probability to each extensional strategy over M_{C_i}.

Running the sheaf theory script

We can now follow the same structure as in Abramsky-Brandenburger, replacing the “flat” event sheaf of local sections by the sheaf of strategies.

We recall the distribution monad $\mathcal{D}_R$, where R is a semiring; when R is the non-negative reals, we recover discrete probability distributions. We have the presheaf $\mathcal{D}_R\Gamma$, obtained by composing the endofunctor part of the monad with Γ.

An **empirical model** over the scenario $(M, \mathcal{C})$ is a family $\{e_i\}$, where $e_i \in \mathcal{D}_R\Gamma(C_i)$, subject to the usual compatibility conditions: for all i, j, $e_i|_{C_i \cap C_j} = e_j|_{C_i \cap C_j}$. Thus e_i assigns a probability to each extensional strategy over M_{C_i}.

The model is **causally non-contextual** if there is a distribution $d \in \mathcal{D}_R\Gamma(X)$ such that, for all i, $d|_{C_i} = e_i$.

Running the sheaf theory script

We can now follow the same structure as in Abramsky-Brandenburger, replacing the "flat" event sheaf of local sections by the sheaf of strategies.

We recall the distribution monad $\mathcal{D}_R$, where R is a semiring; when R is the non-negative reals, we recover discrete probability distributions. We have the presheaf $\mathcal{D}_R\Gamma$, obtained by composing the endofunctor part of the monad with Γ.

An **empirical model** over the scenario $(M, \mathcal{C})$ is a family $\{e_i\}$, where $e_i \in \mathcal{D}_R\Gamma(C_i)$, subject to the usual compatibility conditions: for all i, j, $e_i|_{C_i \cap C_j} = e_j|_{C_i \cap C_j}$. Thus e_i assigns a probability to each extensional strategy over M_{C_i}.

The model is **causally non-contextual** if there is a distribution $d \in \mathcal{D}_R\Gamma(X)$ such that, for all i, $d|_{C_i} = e_i$.

We can show that this recovers

- Standard "flat" contextuality when the enabling is trivial (all measurements initially enabled)
- The Gogioso-Pinzani theory of contextuality for causal Bell scenarios

Experimenter strategies and adaptive computation

But this is only part of the picture!

Experimenter strategies and adaptive computation

But this is only part of the picture!

The strategies considered so far have been strategies for Nature, which choose an outcome for each measurement which can be chosen by the Experimenter.
Using the duality inherent in game theory, there is also a notion of **strategy for Experimenter**.

Experimenter strategies and adaptive computation

But this is only part of the picture!

The strategies considered so far have been strategies for Nature, which choose an outcome for each measurement which can be chosen by the Experimenter.
Using the duality inherent in game theory, there is also a notion of **strategy for Experimenter**.

We define a **strategy for Experimenter** over M to be a set $\tau \subseteq \mathcal{H}(M)$ which is co-total: if s is a non-maximal history in τ, then there is x such that $s \cup \{(x, o)\} \in \tau$ for all $o \in O_x$.

Experimenter strategies and adaptive computation

But this is only part of the picture!

The strategies considered so far have been strategies for Nature, which choose an outcome for each measurement which can be chosen by the Experimenter.
Using the duality inherent in game theory, there is also a notion of **strategy for Experimenter**.

We define a **strategy for Experimenter** over M to be a set $\tau \subseteq \mathcal{H}(M)$ which is co-total: if s is a non-maximal history in τ, then there is x such that $s \cup \{(x, o)\} \in \tau$ for all $o \in O_x$.

Thus at each stage, the Experimenter chooses the next measurement to be performed. It must then accept any possible response from Nature. The future choices of the Experimenter can then depend on Nature's responses, allowing for adaptive protocols.

Experimenter strategies and adaptive computation

But this is only part of the picture!

The strategies considered so far have been strategies for Nature, which choose an outcome for each measurement which can be chosen by the Experimenter.
Using the duality inherent in game theory, there is also a notion of **strategy for Experimenter**.

We define a **strategy for Experimenter** over M to be a set $\tau \subseteq \mathcal{H}(M)$ which is co-total: if s is a non-maximal history in τ, then there is x such that $s \cup \{(x, o)\} \in \tau$ for all $o \in O_x$.

Thus at each stage, the Experimenter chooses the next measurement to be performed. It must then accept any possible response from Nature. The future choices of the Experimenter can then depend on Nature's responses, allowing for adaptive protocols.

We can use Experimenter strategies to capture adaptive MBQC.

The Big Picture

We shall refer to strategies for Nature as N-strategies, and to strategies for Experimenter as E-strategies.

The Big Picture

We shall refer to strategies for Nature as N-strategies, and to strategies for Experimenter as E-strategies.

If we are given an N-strategy σ and an E-strategy τ, we can play them off against each other:

$$\langle \sigma \mid \tau \rangle := \sigma \cap \tau.$$

If τ is deterministic, at each stage τ chooses a unique measurement, and σ a unique outcome for that measurement, so this will be the down-set of a unique maximal history s. In general, it determines a set of histories.

The Big Picture

We shall refer to strategies for Nature as N-strategies, and to strategies for Experimenter as E-strategies.

If we are given an N-strategy σ and an E-strategy τ, we can play them off against each other:

$$\langle \sigma \mid \tau \rangle := \sigma \cap \tau.$$

If τ is deterministic, at each stage τ chooses a unique measurement, and σ a unique outcome for that measurement, so this will be the down-set of a unique maximal history s. In general, it determines a set of histories.

A general causal empirical model will specify a distribution on N-strategies and a distribution on E-strategies for each context.
These distributions can be pushed forward through the evaluation map to yield distributions on histories.

The Big Picture

We shall refer to strategies for Nature as N-strategies, and to strategies for Experimenter as E-strategies.

If we are given an N-strategy σ and an E-strategy τ, we can play them off against each other:

$$\langle \sigma \mid \tau \rangle := \sigma \cap \tau.$$

If τ is deterministic, at each stage τ chooses a unique measurement, and σ a unique outcome for that measurement, so this will be the down-set of a unique maximal history s. In general, it determines a set of histories.

A general causal empirical model will specify a distribution on N-strategies and a distribution on E-strategies for each context.
These distributions can be pushed forward through the evaluation map to yield distributions on histories.

This provides a basis for exploring a wide range of phenomena.

The Big Picture

We shall refer to strategies for Nature as N-strategies, and to strategies for Experimenter as E-strategies.

If we are given an N-strategy σ and an E-strategy τ, we can play them off against each other:

$$\langle \sigma \mid \tau \rangle := \sigma \cap \tau.$$

If τ is deterministic, at each stage τ chooses a unique measurement, and σ a unique outcome for that measurement, so this will be the down-set of a unique maximal history s. In general, it determines a set of histories.

A general causal empirical model will specify a distribution on N-strategies and a distribution on E-strategies for each context.
These distributions can be pushed forward through the evaluation map to yield distributions on histories.

This provides a basis for exploring a wide range of phenomena.

To be continued!